Information-Theoretic Bounds for Adaptive Sparse Recovery

Cem Aksoylar
Department of Electrical and Computer Engineering,
Boston University, Boston, MA 02215

Venkatesh Saligrama
Department of Electrical and Computer Engineering,
Boston University, Boston, MA 02215

arXiv:1402.5731v1 [cs.IT] 24 Feb 2014

***Abstract*—We derive an information-theoretic lower bound for sample complexity in sparse recovery problems where inputs can be chosen sequentially and adaptively. This lower bound is in terms of a simple mutual information expression and unifies many different linear and nonlinear observation models. Using this formula we derive bounds for adaptive compressive sensing (CS), group testing and 1-bit CS problems. We show that adaptivity cannot decrease sample complexity in group testing, 1-bit CS and CS with linear sparsity. In contrast, we show there might be mild performance gains for CS in the sublinear regime. Our unified analysis also allows characterization of gains due to adaptivity from a wider perspective on sparse problems.**

I. Introduction

For many sparse recovery applications, it has been shown that adaptive methods with sequential and flexible measurement designs improve practical performance compared to nonadaptive methods. From a theoretical point of view, adaptive methods should perform at least as well as nonadaptive methods asymptotically, as the latter is a special case of the former. However, it is an interesting problem to determine whether they can perform strictly better for different recovery problems and problem conditions. While such methods have been theoretically analyzed for specific problems of interest, it is not clear at a high level what properties of sparse problems allow adaptive methods to have strictly better recovery performance compared to nonadaptive ones.

In this work we consider a high-level unifying framework and obtain a lower bound on the sample complexity of adaptive sparse recovery problems. Our framework characterizes the problems of interest with the following sparsity assumption: Let $X = (X_1, \dots, X_N)$ be a set of variables and Y be a corresponding observation generated by an observation model $P(Y|X)$ which satisfies the conditional independence property

$$P(Y|X) = P(Y|\{X_k\}_{k\in S}), \tag{1}$$

for some set of variables $S \subset \{1, \dots, N\}$. Then, the aim is to determine the salient set S given T samples of variable/observation pairs, (X^T, Y^T). These pairs are generated sequentially and variables X can be chosen *adaptively* depending on past variables and observations. This framework encapsulates many sparse problems of interest, e.g. support recovery in compressive sensing (CS) [1], its nonlinear extensions (e.g. quantized CS [2]) and other nonlinear problems such as group testing [3].

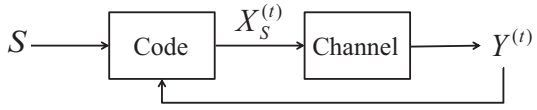

Fig. 1. Channel model representation of adaptive sparse recovery.

As the result of our analysis, we obtain a mutual information formula for the lower bound, which depends on the observation model $P(Y|X_S)$ and the distribution of X, $p_t(X)$, at each step t of the sequence. We obtain this bound using a Fano's inequality type argument, inspired by the proof of the upper bound on capacity for channel coding with feedback [4]. Our result is unifying for all adaptive sparse recovery problems, similar to the results for [5] for the nonadaptive case where variables are generated by a distribution $p(X)$ IID over T samples.

We then obtain adaptive lower bounds for linear and nonlinear applications. We look at the highly nonlinear problem of group testing, where we show that $T = \Omega(K \log(N/K))$ tests are necessary for the adaptive case. This bound can be achieved by nonadaptive methods [3], [5], implying adaptivity cannot improve performance. Similarly, we consider 1-bit CS as a nonlinear extension of linear CS. We again show $T = \Omega(K \log(N/K))$ is a lower bound to argue adaptivity cannot help, as it is matched by nonadaptive upper bounds [6] for sufficiently high SNR. The same phenomenon happens for linear CS with linear sparsity $K = \Theta(N)$. In this case we show SNR $= \Omega(\log N)$ is necessary and the lower bound $T = \Omega(N)$ is achieved by nonadaptive methods for this SNR [7]. However for sublinear sparsity, we show there might be mild gains for T and SNR, consistent with the results of [8], [9].

There is a large body of work on both adaptive recovery methods and lower bounds, however these analyses are fragmented compared to our unifying approach as they only consider specific problems. The linear problem of adaptive CS has been especially well-studied: Lower bounds have been derived for support recovery [10], [11] and adaptive algorithms have been analyzed to obtain upper bounds [8], [9]. There is relatively little work on adaptive recovery on nonlinear models. Adaptive group testing has been investigated by [12], [13]

where lower bounds are derived and adaptive algorithms are analyzed (see ref.s in [13]). Adaptive 1-bit CS algorithms have been proposed [2], however adaptive lower bounds have not been studied to the extent of our knowledge.

The generality of our analysis also allows us to look at the big picture and comment on the role of adaptivity and how it is related to the nature of a problem. We conjecture that adaptivity may help only if there are "sum-power"-like constraints on the variables as in linear CS. If the variables are not constrained and the difficulty of the problem only stems from the observation model, adaptivity does not increase asymptotic performance.

II. An Information-Theoretic Framework for Sparse Recovery

We assume that an observation Y is generated by an observation model for which we assume $P(Y|X) = P(Y|X_S)$, for $S \subset \{1, \dots, N\}$ with $|S| = K$ and $X_S = \{X_k\}_{k \in S}$. We consider the scenario where a latent observation model parameter β_S may exist, with corresponding $P(\beta_S)$ and $P(Y|X_S, \beta_S)$. Finally, let ω index all sets of size K among N variables, such that $\omega \in \left\{1, \dots, \binom{N}{K}\right\}$ and the corresponding set is S_ω. In addition to conditional independence, the only other assumption we make is that ω is chosen uniformly at random among $\binom{N}{K}$ sets.

A simple example of the type of problems we consider is the CS model [1], where the observations are given by $Y = \langle X, \beta \rangle + W$ for a K-sparse vector β with support S, support coefficients β_S and noise W. Another example is the group testing model [3], where X is a Boolean test inclusion vector that determines whether an item is included in the test or not and S is the set of defective items. The group testing model assumes that the test outcomes Y are only dependent on the inclusion of defective items, given by X_S. Further examples and details can be found in [14].

We observe a sequence of T variable-observation pairs $(X^T, Y^T) = (X^{(t)}, Y^{(t)})_{t=1,\dots,T}$. A decoder $g(X^T, Y^T)$ outputs an estimate $\hat{\omega}$ of index ω and we aim to characterize the error probability P_e that $\hat{\omega} \neq \omega$ to obtain conditions on T for successful recovery, in terms of K, N and other problem parameters.

We now present the bounds on sample complexity from [5] and [14]. The following result is a lower bound on the number of samples for recovery, which is the nonadaptive analogue of our main result, Theorem III.1.

Theorem II.1 (Nonadaptive lower bound)**.** *Let $X^T = (X^{(1)}, \dots, X^{(T)})$ be generated IID across $t = 1, \dots, T$ according to $p(X)$. Then, a lower bound (or a necessary condition) on the number of samples required for P_e to be asymptotically bounded away from zero is given by*

$$T \geq \max_{\tilde{S} \subset S_\omega} \frac{\log \binom{N-K}{K-|\tilde{S}|}}{I_{\tilde{S}}}, \tag{2}$$

where $\tilde{S}$ is a proper subset of S_ω and we define

$$I_{\tilde{S}} = I(X_{S_\omega \setminus \tilde{S}}; Y | X_{\tilde{S}}, \beta_{S_\omega}).$$

An upper bound on the number of samples is also presented in [5] and [14]. For the upper bound, it is further assumed that the variables are generated IID across both samples $t = 1, \dots, T$ and variables $n = 1, \dots, N$. The error probability of a Maximum Likelihood decoder is analyzed to obtain a sufficient condition for recovery. The lower bound given in Theorem II.1 is *order-wise* tight as it matches the upper bound, when restricted to IID probability distributions on X^T, for K not scaling with N and provided that a mild condition on the mutual information is satisfied.

III. A Lower Bound for Adaptive Recovery

In this section we analyze the adaptive scenario where at each $t = 1, \dots, T$, $X^{(t)}$ is given by a (possibly random) function $X^{(t)} = f_t(X^{(1:t-1)}, Y^{(1:t-1)})$, where we use $a:b$ to denote the sequence of integers $\{k \in \mathbb{N} : a \leq k \leq b\}$. We state a lower bound on the number of samples in the adaptive case that holds for any distribution of $X^{(1)}$ and functions f_t. The bound depends on the distributions $p_t(X^{(t)})$ marginalized with respect to other sequence indices.

Theorem III.1 (Adaptive lower bound)**.** *Let $X^T = (X^{(1)}, \dots, X^{(T)})$ be generated such that each $X^{(t)}$ is a (random) function of $X^{(1:t-1)}$ and $Y^{(1:t-1)}$. Then, a lower bound on the number of samples required for P_e to be asymptotically bounded away from zero is given by*

$$T \geq \max_{\tilde{S} \subset S_\omega} \frac{\log \binom{N-|\tilde{S}|}{K-|\tilde{S}|}}{\bar{I}_{\tilde{S}}}, \tag{3}$$

where $\tilde{S}$ is a proper subset and we define

$$\bar{I}_{\tilde{S}} = \frac{1}{T} \sum_{t=1}^{T} I(X^{(t)}_{S_\omega \setminus \tilde{S}}; Y^{(t)} | X^{(t)}_{\tilde{S}}, \beta_{S_\omega}),$$

as the average mutual information over the sequence $t = 1, \dots, T$. Each term in the maximization above is also a lower bound.

In the nonadaptive case, X^T is generated IID across samples $t = 1, \dots, T$, therefore $\bar{I}_{\tilde{S}} = I_{\tilde{S}}$ and the above bound reduces to the bound given in Theorem II.1. This also holds for the adaptive case if the variables are chosen such that $X^{(t)}_{S_\omega}$ is identically distributed across the sequence $t = 1, \dots, T$.

In order to obtain bounds for specific applications in the following section, we use two simple methods: We upper bound $I_{\tilde{S}}$ directly for any $p(X)$ in group testing and 1-bit CS, which leads to a trivial upper bound on $\bar{I}_{\tilde{S}}$. We bound $\bar{I}_{\tilde{S}}$ by $\max_{p(X)} I_{\tilde{S}}$ for linear CS, which also gives an upper bound on $\bar{I}_{\tilde{S}}$.

Proof: Let $\omega \in \left\{1, 2, \dots, \binom{N}{K}\right\}$ be the true index of the salient set such that the condition $P(Y|X) = P(Y|X_{S_\omega})$ is satisfied. Suppose a subset of the true support is revealed, denoted by $\tilde{S} \subset S_\omega$, so that only elements of the set $S_\omega \setminus \tilde{S}$ are left to be identified. Define E as the binary error event that the estimate $\hat{\omega} = g(X^T, Y^T)$ is not equal to ω for a decoder g

and let P_e be the probability of this event. We start by writing the uncertainty of ω given $\tilde{S}$,

$$H(\omega|\tilde{S}) = H(S_\omega \setminus \tilde{S}) = \log \binom{N-|\tilde{S}|}{K-|\tilde{S}|} \tag{4}$$

$$= H(\omega|\hat{\omega},\tilde{S}) + I(\omega;\hat{\omega}|\tilde{S}) \tag{5}$$

$$\leq H(\omega|\hat{\omega},\tilde{S}) + I(\omega;X^T,Y^T|\tilde{S}), \tag{6}$$

where (4) is due to ω being chosen uniformly at random, (5) follows from standard entropic identities and (6) is due to data processing inequality and $\hat{\omega}$ being a function of X^T and Y^T. We analyze the first term and write the following inequality:

$$H(\omega|\hat{\omega},\tilde{S}) = H(E,\omega|\hat{\omega},\tilde{S}) - H(E|\omega,\hat{\omega},\tilde{S}) \tag{7}$$

$$= H(E|\hat{\omega},\tilde{S}) + H(\omega|E,\hat{\omega},\tilde{S}) \tag{8}$$

$$\leq 1 + (1-P_e)\,0 + P_e\,H(\omega|\tilde{S}) \tag{9}$$

$$= 1 + P_e \log \binom{N-|\tilde{S}|}{K-|\tilde{S}|}. \tag{10}$$

The two equalities (7) and (8) again follow from standart identities by noting that the second term in (7) is zero since E is determined completely by ω and $\hat{\omega}$. (9) follows by upper bounding the entropy of a binary variable by 1, expanding the conditional entropy for $E=0$ and $E=1$, noting that $E=0$ implies $\omega=\hat{\omega}$ and removing the conditioning on $\hat{\omega}$ on the last term.

We now look at the second term, $I(\omega;X^T,Y^T|\tilde{S})$. We first note the following:

$$\begin{aligned} I(\omega;X^T,Y^T|\tilde{S}) &\leq I(\omega;X^T,Y^T,\beta_{S_\omega}|\tilde{S}) \\ &= I(\omega;X^T,Y^T|\tilde{S},\beta_{S_\omega}) + I(\omega;\beta_{S_\omega}|\tilde{S}) \\ &= I(\omega;X^T,Y^T|\tilde{S},\beta_{S_\omega}), \end{aligned}$$

where the last equality follows from the independence of ω and β_{S_ω}. For this term, we can then write,

$$\begin{aligned} I(\omega;X^T,Y^T|\tilde{S},\beta_{S_\omega}) &= H(X^T,Y^T|\tilde{S},\beta_{S_\omega}) \\ &\quad - H(X^T,Y^T|\omega,\beta_{S_\omega}) \qquad (11) \\ &= \sum_{t=1}^{T} H(X^{(t)},Y^{(t)}|X^{(1:t-1)},Y^{(1:t-1)},\tilde{S},\beta_{S_\omega}) \\ &\quad - H(X^{(t)},Y^{(t)}|X^{(1:t-1)},Y^{(1:t-1)},\omega,\beta_{S_\omega}) \qquad (12) \\ &= \sum_{t=1}^{T} \Big(H(Y^{(t)}|X^{(t)},\tilde{S},\beta_{S_\omega}) \\ &\quad + H(X^{(t)}|X^{(1:t-1)},Y^{(1:t-1)},\tilde{S},\beta_{S_\omega}) \Big) \\ &\quad - \Big(H(Y^{(t)}|X^{(t)},\omega,\beta_{S_\omega}) \\ &\quad + H(X^{(t)}|X^{(1:t-1)},Y^{(1:t-1)},\omega,\beta_{S_\omega}) \Big) \qquad (13) \\ &\leq \sum_{t=1}^{T} I(X^{(t)}_{S_\omega\setminus\tilde{S}};Y^{(t)}|X^{(t)}_{\tilde{S}},\beta_{S_\omega}) \\ &\quad + I(X^{(t)};S_\omega\setminus\tilde{S}|X^{(1:t-1)},Y^{(1:t-1)},\tilde{S},\beta_{S_\omega}) \qquad (14) \\ &= \sum_{t=1}^{T} I(X^{(t)}_{S_\omega\setminus\tilde{S}};Y^{(t)}|X^{(t)}_{\tilde{S}},\beta_{S_\omega}) \triangleq T\bar{I}_{\tilde{S}}, \qquad (15) \end{aligned}$$

where (12) follows from the chain rule of entropy and (13) follows from the fact that $Y^{(t)}$ is independent of $X^{(t')}$ and $Y^{(t')}$ for $t' \neq t$ given $X^{(t)}$ and β_{S_ω}. We note that $X^{(t)}_{\tilde{S}}$ is a function of $X^{(t)}$ and $\tilde{S}$, and $Y^{(t)}$ depends only on $X^{(t)}_{\omega}$ given $X^{(t)}$ and ω to obtain (14). Finally, as $X^{(t)}$ only depends on $(X^{(1:t-1)},Y^{(1:t-1)})$, the second term is zero and (15) follows.

Putting together (6), (10) and (15), we can write

$$\log \binom{N-|\tilde{S}|}{K-|\tilde{S}|} \leq 1 + P_e \log \binom{N-|\tilde{S}|}{K-|\tilde{S}|} + T\bar{I}_{\tilde{S}},$$

which leads to

$$P_e \geq 1 - \left(\frac{T\bar{I}_{\tilde{S}} + 1}{\log \binom{N-|\tilde{S}|}{K-|\tilde{S}|}} \right),$$

and thus for P_e not to be strictly positive, we need

$$T \geq \frac{\log \binom{N-|\tilde{S}|}{K-|\tilde{S}|} - 1}{\bar{I}_{\tilde{S}}}.$$

Considering all proper subsets $\tilde{S} \subset S_\omega$, a lower bound on T for recovery is then given by (3). ∎

The proof is fundamentally different from the feedback proof of [4] since we consider extra latent parameters β_{S_ω}, we have the extra overlap terms $X_{\tilde{S}}$ and we explicitly assume variables depend not only on past outputs but also on past inputs.

IV. Applications

In this section we discuss the implications of the adaptive lower bound for some applications. We will first present results for group testing and 1-bit CS; then we will look at the linear CS model.

A. Group Testing

Group testing is the problem of identifying a set of "defective" items from a larger set, where group tests can be performed which results in a positive outcome if and only if a defective item is included in the test group. Formally, we let X^T denote the Boolean test inclusion matrix for N items and T tests and Y^T denote the binary outcomes of T tests, where each test outcome is given by the formula

$$Y^{(t)} = \bigvee_{k\in S} X^{(t)}_k.$$

Theorem IV.1. *$T = \Theta(K\log(N/K))$ is a lower bound on the number of tests to recover the defective set S using adaptive testing.*

Proof: Consider the case $\tilde{S} = \emptyset$. Then note that $\bar{I}_{\tilde{S}} \leq \frac{1}{T}\sum_{t=1}^{T} H(Y^{(t)}) \leq 1$ as Y is binary. Since $\log \binom{N}{K} = \Theta(K\log(N/K))$, it follows from Theorem III.1 that $T = \Theta(K\log(N/K))$ is a lower bound. ∎

The same lower bound was shown for nonadaptive group testing in by the authors in [3], [5] and for adaptive in [12],

[13][1]. Asymptotically matching upper bounds have also been shown for nonadaptive and adaptive testing, see [13]. In fact, the lower bound can be achieved by choosing entries of X^T IID $\sim$ Bernoulli$(1/K)$ for $K = o(N)$ [3]. Therefore we observe that adaptivity cannot improve performance asymptotically in this sparse recovery problem.

B. 1-bit Compressive Sensing

1-bit CS [6] is interesting as the extreme case of quantized CS models which are of practical importance in many real world applications. Mathematically, we have

$$Y^T = Q(X^T\beta + W^T),$$

where X^T is a $T \times N$ sensing matrix with t-th row corresponding to $X^{(t)}$, β is a K-sparse $N \times 1$ vector with support S and W^T is an IID noise vector. $Q(\cdot)$ is a 1-bit quantizer which outputs 1 if the input is nonnegative and 0 otherwise, for each element in the input vector.

Theorem IV.2. *$T = \Theta(K\log(N/K))$ is a lower bound on the number of measurements to recover the support S of β using adaptive measurements.*

The proof is the same as the proof for group testing since $Y^{(t)}$ are binary measurements. Matching upper bounds for noiseless and noisy variants of the problem (with sufficiently high SNR) have been shown by the authors in [6] and [5] using IID Gaussian measurement matrices. Therefore we have shown that support recovery performance in 1-bit CS cannot be increased asymptotically using adaptive measurements in those SNR regimes.

C. Compressive Sensing

We now look at the CS problem with measurement noise. We have the normalized model [7], [5],

$$Y^T = X^T\beta + W^T,$$

where X^T is the $T\times N$ sensing matrix, β is a K-sparse vector of length N with support S and Y^T is the observation vector of length T. W^T is an IID Gaussian noise vector with variance $1/\text{SNR}$. We assume $\beta_k \in \{-1, +1\}$ w.l.o.g. (since we are obtaining a lower bound) and are IID with equal probability for $k \in S$. We constrain the total power of the entries of X^T similar to [8], [9], where we assume $\sum_{n=1}^N E\left[X_n^{(t)2}\right] = NP_t$ and $\sum_{t=1}^T P_t = 1$, to be consistent with [5] and [7]. As a special case, row-wise power constraints can be enforced by setting $P_t = \frac{1}{T}$. Note that this normalized model uses a different convention from the fixed noise variance models used in [8], [9], [10] where bounds on the minimum magnitude of support coefficients are derived instead of SNR.

In order to obtain a valid lower bound for all distributions on X, we optimize $I(X^{(t)}_{S_\omega\setminus\tilde{S}}; Y^{(t)}|X^{(t)}_{\tilde{S}}, \beta_{S_\omega})$ over probability distributions $p_t(X^{(t)})$, simultaneously for all $\tilde{S} \subset S$. Let $S' = S \setminus \tilde{S}$ and note that

$$\begin{aligned} &I(X^{(t)}_{S_\omega\setminus\tilde{S}}; Y^{(t)}|X^{(t)}_{\tilde{S}}, \beta_{S_\omega}) \\ &\quad= h(Y^{(t)}|X^{(t)}_{\tilde{S}}, \beta_S) - h(Y^{(t)}|X^{(t)}_S, \beta_S) \\ &\quad= h(X^{(t)\top}_{S'}\beta_{S'} + W^{(t)}|X^{(t)}_{\tilde{S}}, \beta_{S'}) - h(W^{(t)}). \end{aligned}$$

Lemma IV.1. *$p_t(X^{(t)})$ that maximizes $\min_{|\tilde{S}|} I(X^{(t)}_{S_\omega\setminus\tilde{S}}; Y^{(t)}|X^{(t)}_{\tilde{S}}, \beta_{S_\omega})$ subject to power constraints sets $\{X^{(t)}_k\}_{k\notin S} = 0$ and is jointly Gaussian in $X^{(t)}_S$, with zero mean and covariance matrix $\Sigma^{(t)}$ with diagonals $\frac{NP_t}{K}$ and the off-diagonals equal to some scalar ρ.*

We do not prove the lemma due to space constraints, however it follows from the fact that β_S is IID and symmetric around zero, the maximization is over all subsets $\tilde{S}$ in (3) and since an entropy is being maximized subject to power constraints.

Then, removing the conditioning on $X_{\tilde{S}}$, $I(X^{(t)}_{S_\omega\setminus\tilde{S}}; Y^{(t)}|X^{(t)}_{\tilde{S}}, \beta_{S_\omega})$ can be upper bounded by

$$\begin{aligned} I(X^{(t)}_{S_\omega\setminus\tilde{S}}; Y^{(t)}|X^{(t)}_{\tilde{S}}, \beta_{S_\omega}) &\le \frac{1}{2}E_{\beta_{S'}}\left[\ln\left(1 + \text{SNR}\beta^\top_{S'}\Sigma^{(t)}_{S'}\beta_{S'}\right)\right] \\ &\le \frac{1}{2}\ln\left(1 + E\left[\beta^\top_{S'}\Sigma^{(t)}_{S'}\beta_{S'}\right]\right), \end{aligned}$$

due to Jensen's inequality, where $\Sigma_{S'}$ is the submatrix of Σ corresponding to the indices in S'. Since $\Sigma_{S'}$ is a circulant matrix, we can write $\Sigma_{S'} = F\Lambda F^\star$ where F is the unitary DFT matrix and $F^\star$ is its conjugate transpose. Λ is a diagonal matrix with its first element equal to $\frac{NP_t}{K} + (|S'| - 1)\rho$ and rest of the diagonals $\frac{NP_t}{K} - \rho$. It is also easy to show that $\widetilde{\beta_{S'}} = F^\star\beta_{S'}$ is also IID and has variances equal to 1. Then it follows that, independent of the value of ρ, we have

$$I(X^{(t)}_{S_\omega\setminus\tilde{S}}; Y^{(t)}|X^{(t)}_{\tilde{S}}, \beta_{S_\omega}) \le \frac{1}{2}\ln\left(1 + \text{SNR}\frac{(K - |\tilde{S}|)NP_t}{K}\right).$$

For $\bar{I}_{\tilde{S}}$ we can then write

$$\begin{aligned} \bar{I}_{\tilde{S}} &\le \frac{1}{T}\sum_{t=1}^T \frac{1}{2}\ln\left(1 + \text{SNR}\frac{(K - |\tilde{S}|)NP_t}{K}\right) \\ &\le \frac{1}{2}\ln\left(1 + \text{SNR}\frac{(K - |\tilde{S}|)N}{KT}\right), \end{aligned}$$

where the second inequality follows from the fact that the sum is maximized by $P_t = \frac{1}{T}$ for all t subject to the constraint $\sum_t P_t = 1$. Note that this implies distributing equal power to all T rows and the same bound holds for the case where row-wise power is constrained (instead of total power of all entries of X^T).

We then evaluate (3) with the above bound on $\bar{I}_{\tilde{S}}$ for (T, SNR) pairs for different $|\tilde{S}|$ to determine which pairs violate the lower bound and obtain the following theorem.

[1][12] inadvertently makes a strong assumption that the variables are identically distributed over the sequence, which is not necessarily true for adaptive testing and leads to the identity $\bar{I}_{\tilde{S}} = I_{\tilde{S}}$. In subsequent work [13] the authors obtain the same lower bound through a different argument.

Theorem IV.3. *A necessary condition on a (T, SNR) pair for exact support recovery of β with adaptive measurements and an arbitrarily small error probability is $T \log\left(1 + SNR\frac{iN}{KT}\right) = \Omega(i \log(N/i))$ for all $i \in \{1, \ldots, K\}$.*

Below corollary then follows by noting that the left-hand side is increasing in T and letting T grow to infinity independent of other quantities for $i = 1$.

Corollary IV.1. *$SNR = \Omega\left(\frac{K \log N}{N}\right)$ is a necessary condition for exact support recovery with adaptive measurements and an arbitrarily small error probability.*

Our result is unique since the lower bound on SNR and the relationship between different T and SNR can be explicitly characterized. Our result proves that for the linear sparsity regime $K = \Theta(N)$, SNR $= \Theta(\log N)$ is necessary and $T = \Theta(N)$ is necessary for that SNR. This shows that adaptivity does not help compared to nonadaptive since SNR $= \Theta(\log N)$ is necessary and $T = \Theta(N)$ is achievable in that case [7].

Corollary IV.1 implies that the necessary condition for adaptive measurements is possibly more relaxed for $K = o(N)$, in fact it is comparable to the best known adaptive lower bound $\Omega(\log K)$ [8]. Similarly, Theorem IV.3 implies that it might be possible to recover S with less measurements with the same noise levels for $K = o(N)$; an example is $T = \Theta(K)$ and SNR $= \Omega(\log N)$, compared to $T = \Theta(K \log(N/K))$ lower bound for nonadaptive recovery with the same SNR; or $T = \Theta(K \log N)$ for SNR $= \Omega(\log K)$, which is achievable for $K = o(N)$ [8]. Therefore, in contrast to the group testing and 1-bit CS examples, in linear CS we have shown that there might be room for improvement using adaptive measurements in the sublinear sparsity regime. This result is consistent with previous lower bounds for adaptive CS [9], [11].

Finally, note that performance improvements would not be possible if we constrained the power of each element of the sensing matrix individually or constrained the total power of $\langle X, \beta \rangle$, as in both cases it would not be possible to improve measurements by "concentrating" power on X_S with the information from previous measurements.

V. Discussion

Looking at Theorem III.1 and the applications discussed above, we can argue that sample complexity is minimized in cases where the average information in the sequence $\bar{I}_{\tilde{S}}$ can approach to the maximum value (over distributions on X) of $I_{\tilde{S}}$. However, in models where the only difficulty is the uncertainty in the observation model $P(Y|X_{S_\omega})$ and there are no total-power restrictions on measurements X (even if there are element-wise restrictions), an optimal $p(X)$ can be determined beforehand. This $p(X)$ would maximize $I_{\tilde{S}}$ for all possible sets ω simultaneously, eliminating the need for adaptive measurements. Group testing is an example of such a problem, since there are no restrictions on the Boolean testing matrix and the only uncertainty is due to the observation model, and hence we see no gains from adaptivity asymptotically.

However, in the case of a total-power constraint (such as the linear CS model), adaptivity helps by obtaining more information on ω and transferring power from less likely candidates for S_ω to more likely ones, therefore increasing the "effective SNR" as the measurement sequence progresses. Since nonadaptive measurements do not have any prior information on ω, they attempt to distribute power evenly over all N candidate indices so they cannot achieve the same performance asymptotically, at least for sublinear sparsity.

1-bit CS is an interesting example since SNR and the total power of measurement vectors are still important but we observe that for sufficiently high SNR the binary output constrains performance more than the measurement power. Therefore nonadaptive methods can achieve the adaptive lower bound asymptotically. However, it is an open question whether adaptivity can increase asymptotic performance for lower SNR values similar to linear CS.

Acknowledgements

This work was supported by NSF Grant 0932114 and NSF Grant CCF-1320547.

References

[1] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289 –1306, April 2006.

[2] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," in *Proc. of Conf. on Information Sciences and Systems (CISS)*, March 2008, pp. 16–21.

[3] G. Atia and V. Saligrama, "Boolean compressed sensing and noisy group testing," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, March 2012.

[4] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: John Wiley and Sons, Inc., 1991.

[5] C. Aksoylar, G. Atia, and V. Saligrama, "Sparse signal processing with linear and non-linear observations: A unified Shannon theoretic approach," in *Proc. of the IEEE Inf. Theory Workshop (ITW)*, Seville, Spain, 2013.

[6] A. Gupta, R. Nowak, and B. Recht, "Sample complexity for 1-bit compressed sensing and sparse classification," in *Proc. of the IEEE Int. Symp. on Inf. Theory (ISIT)*, June 2010, pp. 1553–1557.

[7] S. Aeron, M. Zhao, and V. Saligrama, "Information theoretic bounds for compressed sensing," *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 5111–5130, Oct. 2010.

[8] M. L. Malloy and R. D. Nowak, "Near-optimal adaptive compressed sensing," in *Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*. IEEE, 2012, pp. 1935–1939.

[9] J. D. Haupt, R. G. Baraniuk, R. M. Castro, and R. D. Nowak, "Compressive distilled sensing: Sparse recovery using adaptivity in compressive measurements," in *Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*. IEEE, 2009, pp. 1551–1555.

[10] E. Arias-Castro, E. Candes, and M. Davenport, "On the fundamental limits of adaptive sensing," *IEEE Trans. Inf. Theory*, vol. 59, no. 1, pp. 472–481, 2013.

[11] R. M. Castro, "Adaptive Sensing Performance Lower Bounds for Sparse Signal Detection and Support Estimation," *arXiv pre-print, arXiv:1206.0648*, Jun. 2012.

[12] M. Aldridge, "Adaptive group testing as channel coding with feedback," in *Proc. of the IEEE Int. Symp. on Inf. Theory (ISIT)*, 2012, pp. 1832–1836.

[13] L. Baldassini, O. Johnson, and M. Aldridge, "The capacity of adaptive group testing," in *Proc. of the IEEE Int. Symp. on Inf. Theory (ISIT)*, 2013, pp. 2676–2680.

[14] C. Aksoylar, G. Atia, and V. Saligrama, "Sparse Signal Processing with Linear and Non-Linear Observations: A Unified Shannon Theoretic Approach," *arXiv pre-print, arXiv:1304.0682*, Apr. 2013.